

PCT


WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁷ : G06F 12/08		A1	(11) International Publication Number: WO 00/34870
			(43) International Publication Date: 15 June 2000 (15.06.00)
(21) International Application Number: PCT/US99/21249		(81) Designated States: AE, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CR, CU, CZ, DE, DK, DM, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).	
(22) International Filing Date: 22 September 1999 (22.09.99)			
(30) Priority Data: 09/207,278 8 December 1998 (08.12.98) US			
(63) Related by Continuation (CON) or Continuation-in-Part (CIP) to Earlier Application US 09/207,278 (CON) Filed on 8 December 1998 (08.12.98)			
(71) Applicant (for all designated States except US): INTEL CORPORATION [US/US]; 2200 Mission College Boulevard, Santa Clara, CA 95052 (US).		Published With international search report.	
(72) Inventors; and (75) Inventors/Applicants (for US only): CAI, Zhong-ning [CN/US]; 3219 N.W. Chapin Drive, Portland, OR 97229 (US). NAKANISHI, Tosaku [JP/US]; 2048 Walnut Avenue, Saratoga, CA 95070 (US).			
(74) Agent: TROP, Timothy; Trop, Pruner, Hu & Miles, P.C., Suite 100, 8554 Katy Freeway, Houston, TX 77024 (US).			

(54) Title: BUFFER MEMORY MANAGEMENT IN A SYSTEM HAVING MULTIPLE EXECUTION ENTITIES



(57) Abstract

A system includes multiple program execution entities (e.g., tasks, processes, threads, and the like) (258, 260) and a cache memory (100) having multiple sections (102, 104, 106). An identifier is assigned to each execution entity. An instruction of one of the execution entities (258, 260) is retrieved and an associated identifier is decoded. Informations associated with the instruction is stored in one of the cache sections (102, 104, 106) based on the identifier.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

- 1 -

Buffer Memory Management In A System Having
Multiple Execution Entities

Background

The invention relates to buffer memory management in a system having multiple execution entities.

A buffer memory can be a relatively small, fast memory placed between a memory
5 and another device that is capable of accessing the memory. An example of a buffer memory is a cache memory located between a processor and system memory (which typically is relatively large and slow) to reduce the effective access time required by the processor to retrieve information from the system memory. In some systems, a multi-level cache system may be used for further performance improvement. A first-level cache (L1 cache) may be
10 implemented in the processor itself, and a second-level, typically larger cache (L2 cache) is externally coupled to the processor.

Further, in some conventional memory systems, a cache memory may include separate instruction and data cache units, one to store instructions and the other to store data. During operation, a processor may fetch instructions from system memory to store in the
15 instruction cache unit. Data processed by those instructions may be stored in the data cache unit. If information, such as instruction or data, requested by the processor is already stored in cache memory, then a cache memory hit is said to have occurred. A cache memory hit reduces the time needed for the processor to access information stored in memory, which improves processor performance.

20 However, if information needed by the processor is not stored in cache memory, then a cache miss is said to have occurred. When a cache miss occurs, the processor has to access the system memory to retrieve the desired information, which results in a memory access time performance reduction while the processor waits for the slower system memory to respond to the request. To reduce cache misses, different cache management policies have
25 been implemented. One of several mapping schemes may be selected, for example, including a direct mapping scheme or a set associative cache mapping scheme. A set associative cache memory that implements k-way associative mapping, e.g., 2-way associative mapping, 4-way associative mapping, and so forth, generally provides a higher hit ratio than direct mapped cache memory. One of several replacement policies may also be specified to improve cache
30 memory hit ratios, including a first-in-first-out (FIFO) or least recently used (LRU) policy.

- 2 -

Another feature of a cache memory that may be configured is the cache memory update policy that specifies how the system memory is updated when a write operation changes the contents of the cache. Update policies include a write-through policy or a write-back policy.

Conventionally, a system, such as a computer, may include multiple application
5 programs and other software layers that have different data flow needs. For example, a program execution entity, such as a process, task, or thread, associated with a multimedia application may transfer large blocks of data (e.g., video data) that are typically not reused. Thus, access of these types of data may cause a cache to fill up with large blocks of data that are not likely to be reused.

10 In filling a cache memory, data used by one execution entity may replace data used by another execution entity, a phenomenon referred to as data cache pollution. Data cache pollution caused by the activities of one execution entity may increase the likelihood of cache misses for another execution entity, which may reduce overall system performance.

A need thus exists for a memory architecture that provides improved memory
15 performance.

Summary

In general, according to an embodiment, a system includes a processor and a plurality of execution entities executable on the processor. A buffer memory in the system has
20 multiple buffer sections. Each buffer section is adapted to store information associated with requests from a corresponding one of the multiple execution entities.

Other features will become apparent from the following description and from the claims.

Brief Description Of The Drawings

25 Fig. 1 is a block diagram of portions of a buffer or cache memory having multiple sections according to an embodiment of the invention.

Fig. 2 is a block diagram of an embodiment of a system including the cache memory of Fig. 1.

30 Fig. 3 illustrates the components of each cache module in the cache memory of Fig. 1.

Fig. 4 is a block diagram of a processor including the cache memory of Fig. 1 along with associated control logic.

- 3 -

Fig. 5 is a flow diagram of an instruction execution sequence performed in the processor of Fig. 4.

Fig. 6 is a flow diagram of an operating system in the system of Fig. 2 that sets up a cache memory according to an embodiment.

5

Detailed Description

In the following description, numerous details are set forth to provide an understanding of the present invention. However, it is to be understood by those skilled in the art that the present invention may be practiced without these details and that numerous variations or modifications from the described embodiments may be possible.

Some embodiments of the invention include a system having a buffer memory that includes several individual buffer sections at one level of the memory hierarchy. Each buffer section may be a separate buffer module or may be a portion of a buffer memory that is separately addressable (that is, memory is separated into different address spaces). The individual buffer sections may be separately configurable and may be assigned to store information of different program execution entities in the system. Such a buffer memory may be referred to as a multi-unit buffer memory.

In some embodiments, the buffer memory may include a cache memory used in any of a variety of applications, e.g., processor subsystems, peripheral device controllers (such as video controllers, hard disk drive controllers, and so forth), and other types of control devices. Systems including such cache memories may include a general-purpose or special-purpose computer, a hand-held electronic device (e.g., telephones, calendar systems, electronic game devices, and the like), appliances, set-top boxes, and other electronics systems. A cache memory having multiple cache sections may be referred to as a multi-unit cache memory. A cache memory section may include a separate cache module or a portion of the cache memory that is separately addressable. The following described embodiments include a computer having a multi-unit cache memory with multiple independent cache modules -- it is to be understood, however, that further embodiments may include computers having multi-unit cache memories with other independently configurable cache sections or other types of systems with buffer memories.

According to some embodiments, the attributes of each individual cache module in a multi-unit cache memory may be independently configurable. Such attributes may include

- 4 -

each cache module's size, organization (e.g., direct mapped versus set associative mapping), replacement policy, update policy, and so forth. Thus, for example, one cache module may be configured to be a direct mapped cache while another cache module may be configured as a k-way set associative cache. The cache modules may also be configured to have different
5 update policies, including a write-through policy or a write-back policy. Other attributes may also be set differently for the different cache modules, as further described below.

Some processors may be capable of receiving requests from multiple execution entities for processing. A processor may include, by way of example, a general-purpose or a special-purpose microprocessor, a microcontroller, or other types of control devices such as
10 application-specific integrated circuits (ASICs), programmable gate arrays (PGAs), and the like. A program execution entity according to one embodiment may be the basic unit of work of software and firmware layers that are loaded in the system. Such basic units of work may include processes, tasks, threads, or other units, as definable according to different systems. For example, in some operating systems, such as certain Windows® operating systems by
15 Microsoft Corporation, multiple threads associated with processes in the system may be executable by the processor to perform different operations. Another operating system that offers multithreading or multitasking capabilities is the Be Operating System (BeOS) from BE, Inc., as described in the BE Operating System Product Data Sheet, published at <http://www.be.com> in 1998.

20 In such operating systems, multiple execution entities associated with different software and firmware layers may be active at a time. Requests from these execution entities are scheduled by the operating system according to a predetermined priority protocol, e.g., round-robin, etc. Such operating systems are said to be multitasking or multithreading operating systems. To take advantage of the multitasking or multithreading capabilities of a
25 system, the independent cache modules of a multi-unit cache memory may be assigned to store information of corresponding execution entities. Thus, for example, execution entities of a multimedia application may be assigned to one cache module, while execution entities of other applications may be assigned to different cache modules of the multi-unit cache memory. To that end, according to one embodiment, requests from each execution entity
30 may be assigned to different execution entity identifiers (EIDs). Thus, requests from execution entities of a first application may be assigned to one EID, and requests from another execution entity may be assigned another EID. Thus, according to this embodiment,

- 5 -

a cache module may be configured for the general data usage behavior of an assigned application.

In another embodiment, the execution entities created by one software or firmware layer may further be subdivided to have multiple EIDs. For example, an application may
5 create execution entities that process data according to different temporal and spatial locality characteristics. For example, some execution entities may be more likely to reuse data than other execution entities created by the same application. Thus, it may be beneficial to further separately assign these different execution entities to different cache modules in the multi-unit cache memory. Thus, in one alternative embodiment, requests from different execution
10 entities of one application may be assigned more than one EID so that different cache modules may be utilized. In addition, execution entities of different applications may be assigned the same EID. Thus, for example, a first execution entity of a multimedia application may be assigned EID 1, while a second execution entity of the multimedia application may be assigned EID 2. In the same system, execution entities of a spreadsheet
15 application having similar data usage characteristics as the second execution entity of the multimedia application may also be assigned EID 2.

In further embodiments, other different schemes may be implemented in assigning EIDs to requests of execution entities. Based on the EID associated with an instruction, a cache controller for the cache memory can keep track of which cache module of the multi-unit cache memory is to be used to store data accessed by the instruction. As a result, cache
20 utilization may be improved since the individual cache modules may be configured to take advantage of the data usage characteristics of the different execution entities in the system. For example, a multimedia application may typically generate requests that transfer large blocks of data that are not re-used. A cache module assigned to these types of requests may
25 be configured to implement the FIFO replacement policy and write-through update policy. Cache modules assigned to other types of requests may have different configurations.

As execution entities are created in a system, EID identifiers may be assigned to these execution entities by an operating system. Referring to Fig. 6, according to one embodiment, if a new execution entity is detected (at 502), the operating system may access (at 504)
30 configuration information loaded during system initialization to determine how EID identifiers are to be assigned. The operating system next assigns (at 506) the appropriate EID identifier to the execution entity. For example, the operating system may be able to assign

- 6 -

three EIDs to correspond to three cache modules in a multi-unit cache memory. Execution entities having one general data usage characteristic may be assigned a first EID identifier, and execution entities having a second general data usage characteristic may be assigned a second EID identifier. A default EID identifier may be assigned to those execution entities
5 that are not specifically assigned one of the other two EID identifiers.

In addition, based on the configuration information, the operating system also assigns (at 508) certain attributes of each cache module in the multi-unit cache memory. Such attributes may include the update, replacement, and placement policies. The operating system may also assign the attributes for the default cache module of the multi-unit cache
10 memory. In alternative embodiments, the EID identifiers and cache attributes may be assigned as described above by a software layer that is separate from the operating system.

In one example, execution entities of a multimedia application that transfer large amounts of data and that do not typically reuse the data may be allocated one EID identifier so that such data are stored in a first cache module configured for the cache data usage
15 characteristics of these execution entities. Execution entities of intensive arithmetic applications, such as compression applications, may be assigned another EID identifier so that data is stored in another cache module that is configured for cache data operations characterized by increased spatial locality.

In some embodiments, a multi-unit cache memory having multiple cache modules
20 may be implemented in a multilevel cache memory having multiple levels of cache memory (e.g., an L1 cache and an L2 cache). Such a cache memory may be referred to as a multilevel, multi-unit cache memory, in which at least one level includes a multi-unit cache memory. Thus, for example, a multilevel, multi-unit cache memory having two levels may be constructed in the following manner: the first level is a multi-unit cache and the second
25 level is a conventional cache; the first level is a multi-unit cache and the second level is a multi-unit cache; or the first level is a conventional cache and the second level is a multi-unit cache.

The individual cache modules of a multi-unit cache may be referred to as P-caches. Thus, for example, a multi-unit cache memory may include several P-caches, including a
30 P0-cache, a P1-cache, a P2-cache, and so forth. The different P-caches may be implemented as separate memory elements or modules, e.g., multiple static random access memory (SRAM) or multiple dynamic random access memory (DRAM) devices. Alternatively,

- 7 -

multiple P-caches may be implemented in one memory device that is sub-divided into separate sections to correspond to the different P-caches. In addition, the multi-unit cache memory may be integrated in another device, e.g., a processor or other control device in a system. Alternatively, the multi-unit cache memory may be a stand-alone unit accessible by
5 control devices to retrieve cached data. In further embodiments, one portion of the multi-unit cache memory may be located in one integrated device while another portion of the multi-unit cache memory is located in another device.

In some embodiments of the invention, each individual P-cache module in a multi-unit cache system may have different attributes, including cache size and organization and
10 cache update, placement, and replacement policies. A placement policy may be specified for each P-cache to determine how information is placed into unfilled portions of the cache. A cache replacement policy is specified to manage replacement of information stored in each P-cache. Example replacement policies may include a first-in-first-out (FIFO) policy, a least-recently-used (LRU) policy, or some other type of replacement policy. A cache update
15 policy manages how information is to be updated when a write occurs to the cache, which may include a write-through policy or a write-back policy.

Referring to Fig. 1, a multi-unit cache memory 100 according to an embodiment includes several P-caches, shown as a P0-cache 102, a P1-cache 104, and a P2-cache 106. A cache controller 108 is associated with the P0-, P1-, and P2-caches 102, 104, and 106. In one
20 embodiment, separate address and data buses may be coupled to each of the cache modules 102, 104 and 106 so that the cache modules may be accessed concurrently. Alternatively, a common address and data bus may be coupled to the cache modules. The cache controller 108 provides control signals to each of the P-cache modules 102-106.

The cache controller 108 includes storage elements 118, in the form of registers or the
25 like, that are programmable by the operating system to specify the EID identifiers associated with each of the P-caches. When the multi-unit cache memory 100 is accessed, the cache controller 108 selects one of the P-caches based on a comparison of the EID provided by a request and the EID values stored in the storage elements 118.

The cache controller 108 also includes a replacement and update control block 120 to
30 control the replacement and update policies of the three separate cache modules, as determined by control information programmed in the storage elements 118. Thus, for

- 8 -

example, the storage elements 118 may be programmed to indicate a FIFO replacement policy for one P-cache and an LRU replacement policy for another P-cache.

The cache controller 108 may also include a tag compare block 122 to compare the tag of an incoming request to the tag stored in the selected one or ones of the P-caches to
5 determine if a cache hit has occurred. Further, if update of main memory 206 (Fig. 2) is needed, a write-back buffer 124 stores the cache line of one of the P-caches to transfer to main memory 206 or an L2 cache 204 (Fig. 2).

To ensure cache data integrity, the cache controller 108 also includes a cache coherency block 126 that determines if an accessed location of a cache module is valid. In
10 one embodiment, each cache module may store a valid/invalid bit. Alternatively, a more sophisticated coherency protocol may be implemented, such as the Modified, Exclusive, Shared, and Invalid (MESI) protocol.

Other control signals that may be provided to the cache controller 108 may include a cache disable (CD) signal and a cache flush (CF) signal. In addition, other cache related
15 signals such as snoop signals may be provided to the cache controller 108.

Referring to Figure 2, the multi-unit cache system 100 may be implemented in a number of different locations (e.g., processor subsystem, bridge controllers, peripheral device controllers, storage device controllers, and the like) in a system 10. In one embodiment, the system 10 includes a computer, although in alternative embodiments, the system 10 may be
20 any other electronic device in which a cache or buffer memory may be implemented.

The system 10 includes a central processing unit (CPU) 200, which may include a processor or other suitable control device, having one or more levels of cache memory. For example, as illustrated, the CPU 200 may include an internal cache that is the level-one (L1) cache 202. In addition, the CPU 200 may be coupled over a host bus 203 to access an
25 external cache that is the level-two (L2) cache 204. The L1 cache 202 may include a code component (for storing instructions) and a data component (for storing data). Similarly, the L2 cache 204 may include code and data components. Thus, instructions and data fetched from main memory 206 are stored in the code and data components, respectively, of the L1 or L2 cache 202 or 204. In other embodiments, separate code and data cache components are
30 not implemented.

In some embodiments, the multi-unit cache memory 100 (Fig. 1) may be implemented in the L1 cache 202, the L2 cache 204, or both. For purposes of this discussion, it is assumed

- 9 -

that the multi-unit cache memory 100 of Fig. 1 is implemented in the L1 cache 202 that is the internal cache of the CPU 200. It is to be understood, however, that the multi-unit cache memory described, or modifications of such a cache memory, may be implemented in the L2 cache 204 or in other controllers in the system, such as a video controller or a hard disk drive controller, as examples. In addition, in this embodiment, the multi-unit cache memory 100
5 forms the data cache component of the L1 cache 202.

The main memory 206 is controlled by a memory controller 207 in a memory hub 208 coupled to the CPU 200 over the host bus 203. In addition, the memory hub 208 may include a cache controller 205 operatively coupled to the L2 cache 204. The memory hub 208 may
10 also include a graphics interface 211 that is coupled over a link 209 to a graphics controller 210, which is in turn coupled to a display 212. As an example, the graphics interface may be according to the Accelerated Graphics Port (A.G.P.) Interface Specification, Revision 2.0, published in May 1998.

The memory hub 208 may also be coupled to an input/output (I/O) hub 214 that
15 includes bridge controllers 215 and 223 coupled to a system bus 216 and a secondary bus 224, respectively. As an example, the system bus may be a Peripheral Component Interconnect (PCI) bus, as defined by the PCI Local Bus Specification, Production Version, Revision 2.1, published in June 1995. The system bus 216 may be coupled to a storage
20 controller 218 that controls access to one or more mass storage devices 220, including a hard disk drive, a compact disc (CD) drive, or a digital video disc (DVD) drive. In an alternative embodiment, the storage controller 218 may be integrated into the I/O hub 214, as may other control functions. The system bus 216 may also be coupled to other components, including, for example, a network controller 222 that is coupled to a network port (not shown).

On the secondary bus 224, additional devices 226 may be coupled, as may be a
25 non-volatile memory 228 that may store power up routines, such as basic input/output system (BIOS) routines. The secondary bus 224 may also include ports for coupling to peripheral devices. Although the description makes reference to specific configurations and architectures of the various layers of the system 10, it is contemplated that numerous modifications and variations of the described and illustrated embodiments may be possible.
30 For example, instead of memory and I/O hubs, a host bridge controller and a system bridge controller may provide equivalent functions, with the host bridge controller coupled between the CPU 200 and the system bus 216 and the system bridge controller 224 coupled between

- 10 -

the system bus 216 and the secondary bus 224. In addition, any of a number of bus protocols may be implemented.

Various different program execution entities are executable by the CPU 200 in the system 10. As illustrated, according to one embodiment, multiple processes 252, 254, and
5 256 are loaded under an operating system 250, which may be a Windows® operating system, for example. Each process may generate one or more execution entities that form the basic units of work in the system. In one example, the execution entities may be threads; as illustrated in Fig. 2, the process 252 may include threads 258 and 260, the process 254 may include a thread 262, and the process 256 may include threads 264 and 266.

10 Various software or firmware (formed of modules, routines, or other layers, for example), including applications, operating system modules or routines, device drivers, BIOS modules or routines, and interrupt handlers, may be stored or otherwise tangibly embodied in one or more storage media in the system. Storage media suitable for tangibly embodying software and firmware instructions may include different forms of memory including
15 semiconductor memory devices such as dynamic or static random access memories, erasable and programmable read-only memories (EPROMs), electrically erasable and programmable read-only memories (EEPROMs), and flash memories; magnetic disks such as fixed, floppy and removable disks; other magnetic media including tape; and optical media such as CD or DVD disks. The instructions stored in the storage media when executed cause the system 10
20 to perform programmed acts.

The software or firmware can be loaded into the system 10 in one of many different ways. For example, instructions or other code segments stored on storage media or transported through a network interface card, modem, or other interface mechanism may be loaded into the system 10 and executed to perform programmed acts. In the loading or
25 transport process, data signals that are embodied as carrier waves (transmitted over telephone lines, network lines, wireless links, cables and the like) may communicate the instructions or code segments to the system 10.

The execution entities (in this case threads) are adapted to perform different operations. For example, a spreadsheet process may create a first thread to perform
30 calculations on entries entered by a user and a second thread to transfer the calculated data into main memory 206. Each thread or execution entity is able to generate requests, which

- 11 -

are stored as instructions in main memory 206. These instructions are fetched by the CPU 200 from main memory 206 for execution.

According to some embodiments, an execution entity identifier (EID) may be assigned to each execution entity running in the system 10. The EID of each execution entity
5 may be assigned by the operating system. In one embodiment, when a scheduler 270 schedules requests from the execution entities for processing by the CPU 200, the associated EID of each execution entity is stored along with one or more corresponding instructions. In this embodiment, the CPU 200 fetches the associated EIDs along with the instructions.

In an alternative embodiment, the EIDs are not stored into memory 206 along with
10 instructions. Instead, multiple instruction memory regions may be defined in the memory 206 to correspond to the different EIDs. Instructions associated with a request from an execution entity having a first EID may be stored in a first instruction memory region; instructions associated with a request from an execution entity having a second EID may be stored in a second instruction memory region; and so forth. In this alternative embodiment,
15 the CPU 200 fetches instructions from memory 206 without associated EIDs. However, based on which of the instruction memory regions the instruction is fetched, the CPU 200 can determine the EID of the instruction.

In yet a further embodiment, in which EIDs are similarly not stored along with instructions, the CPU 200 may include multiple microsequencers assigned to different
20 threads. Thus, one microsequencer may retrieve instructions associated with one thread, another microsequencer may retrieve instructions associated with another thread, and so forth. Each microsequencer may be configured to know locations of instructions of corresponding execution entities. In this embodiment, an instruction's EID may be determined depending on which of the microsequencers fetched that instruction. The
25 determined instruction may then be stored inside the CPU.

The retrieved or determined EID is decoded by the cache controller 108 or by some other suitable decoder to identify which P-cache is to be used when the instruction requests an access to data. The cache controller 108 accesses one of the P-caches to retrieve or store data processed by the corresponding instruction. With the example configuration of Fig. 1,
30 data associated with instructions having EID 0 may be stored in the P0-cache 102, data associated with instructions having EID 1 may be stored in the P1-cache 104, and data associated with instructions having EID 2 may be stored in the P2-cache 106. In some

- 12 -

embodiments, a P-cache may be associated with more than one EID. Further, execution entities from different application and software layers may be assigned the same EID.

Referring to Fig. 3, the general architecture of one of the P-caches is illustrated. In the example shown in Fig. 3, a 4-way set-associative cache is illustrated. Other
5 configurations are also possible, including a direct mapped cache or other k-way set-associative caches. Each P-cache may include a status array 160, a tag array 162, and a data array 164. As illustrated, each of the status array 160, tag array 162, and data array 164 is divided into 4 different sections for the 4-way set-associative organization.

The status array 160 may contain one or more of the following fields: an EID
10 identifier; replacement selection bits (RPS) that are used by the replacement and update control block 120 to replace a cache line; and cache coherency protocol bits. For example, each block of the P-cache module may be associated with a valid/invalid bit to indicate if the corresponding cache location is valid or invalid. Alternatively, the status array 160 may store MESI bits. The replacement selection bits RPS may be used to indicate which cache line is
15 to be replaced. The RPS bits may be used to keep track of the least recently used cache line (for LRU replacement) or the first entered line (for FIFO replacement), as examples.

The cache controller 108 may be implemented as an integrated unit or as several separate control units. As discussed, when an instruction is fetched for execution, the EID associated with the instruction is retrieved. Based on the EID value, the appropriate one of
20 the P-cache modules is selected to retrieve data from or write data to. A hit or miss may be returned depending on whether a valid copy of the associated data is stored in the selected P-cache module.

A multi-unit cache system having independently configurable cache modules according to some embodiments may have one or more of the following advantages. Greater
25 cache management flexibility may be available, since the placement, replacement, and update policies and cache size and organization of each of the P-cache modules may be set to improve cache utilization for corresponding execution entities. Cache performance may be improved by configuring cache modules to take advantage of different cache usage characteristics (to store data or instructions) of different execution entities. Data cache
30 pollution by the different active execution entities in the system 10 may be reduced, which may improve the cache hit ratio. In addition, the multi-unit data cache system may offer high access bandwidth by increasing parallelism for a multithreading or multitasking processor

- 13 -

since the P-cache modules may be concurrently accessible. Such concurrent data cache accesses may help reduce data cache latency to help meet the data access bandwidth demands of high-performance processors.

In another embodiment, compilers for different application programs may
5 dynamically reconfigure attributes of the multi-unit cache memory to further enhance cache performance. For example, during operation, statistical information associated with different execution entities may be collected and stored. Depending on the collected statistical information, the attributes of each P-cache module may be changed. Thus, for example, if a
FIFO replacement policy is determined not to be efficient for a particular P-cache module,
10 the cache controller 108 may be notified to change the replacement policy to the LRU policy or some other replacement policy. This alternative embodiment may provide the flexibility of dynamically changing the configuration of individual P-cache modules in response to how execution entities in the system 10 are performing.

Referring to Fig. 4, in one embodiment, the CPU 200 includes the multi-unit L1 cache
15 memory 202 and associated logic. The multi-unit L1 cache memory includes the three data cache modules: the P0-, P1-, and P2-caches 102, 104, and 106. The P0-cache 102 may be designated as the default data cache that is used to store data associated with execution entities that have not specifically been assigned to one of the other P-caches in the L1 cache. For example, such execution entities may be assigned a default EID 0 by the operating
20 system. The P1- and P2-caches 104 and 106 may be assigned to store data for requests from execution entities having EIDs 1 and 2, respectively. In one embodiment, the P0-cache may be a larger memory than either the P1- or P2-cache since it is the default data cache.

Other components of the CPU 200 according to one example configuration are illustrated in Fig. 4. A bus front unit (BFU) 404 forms the interface to the front side or host
25 bus 203. The BFU 404 may include address drivers and receivers, write buffers, data bus transceivers, bus control logic, bus master control, and parity generation and control.

The instruction path is first described below. Instructions retrieved by the BFU 404 from either the main memory 206 or from the L2 cache 204 may be stored in an instruction cache 406 that is part of the L1 cache 202. The internal instruction cache 406 may keep
30 copies of the most frequently used instructions. According to some embodiments, instructions are fetched along with EIDs from either the main memory 206 or the L2 cache 204 and stored in the instruction cache 406. An instruction buffer and decode logic 408

- 14 -

decodes a selected instruction and associated EID from the instruction cache 406 and produces one or more micro-operations along with corresponding EIDs.

In an alternative embodiment, the instructions are stored in different instruction memory regions of the memory 206 according to different EIDs. However, in this
5 embodiment, the EIDs are not stored along with the instructions. When the CPU 200 fetches an instruction, an associated EID is not retrieved. Instead, the CPU 200 determines the EID of the fetched instruction based on the address location where the instruction is stored. This may be performed, for example, by the decode logic 408. Thus, the EID of an instruction is determined based on which instruction memory region the instruction is fetched from. Once
10 the EID is determined by the CPU 200, it can be attached to the decoded micro-operations and stored in the instruction queue 412.

In yet a further embodiment in which EIDs are not stored along with instructions in memory, multiple program counters and microsequencers may be included in the CPU 200 that are assigned to corresponding threads. This embodiment is described further below.

15 The output port of the instruction buffer and decode logic 408 may be coupled to an instruction queue 412, which stores the micro-operations along with associated EIDs. The output port of the instruction queue 412 is routed to a sequencer 414. The sequencer 414 may include multiple microsequencer units 430, 432, and 434 corresponding to the different EIDs. For example, the microsequencer unit 430 may be configured to process micro-operations
20 associated with EID 0, the microsequencer 432 may be configured to process micro-operations associated with EID 1, and the microsequencer 434 may be configured to process micro-operations associated with EID 2. The micro-operations processed by the microsequencers 430, 432, and 434 are received from the instruction queue 412. According to one embodiment, the microsequencers 430, 432, and 434 may operate simultaneously to
25 process micro-operations associated with different EIDs. Operation of the microsequencers 430, 432, and 434 are controlled by a control logic 436 in the sequencer 414.

In one embodiment, the EIDs associated with micro-operations are originally retrieved from memory along with instructions of the different execution entities. In a further embodiment in which instructions are not stored along with instructions, each
30 microsequencer may be independently configured to fetch instructions of corresponding execution entities. Thus, a first microsequencer fetches instructions associated with a first execution entity, a second microsequencer fetches instructions associated with a second

- 15 -

execution entity, and so forth. The EIDs of a fetched instruction may thus be determined based on which of the microsequencers fetched the instruction.

An execution entity typically includes a number of instructions that are executed in some program order. By default, instruction addresses are simply incremented to fetch the
5 next instruction. If a jump or other conditional branch occurs, then a target address is specified for the address of the next instruction. Thus, the address of the memory location where the next instruction is stored is known. A program counter may be used to keep track of the program order of instructions. A microsequencer works in conjunction with the program counter to execute instructions. To fetch an instruction, the microsequencer may ask
10 a fetch unit, located in the BFU 404 for example, to fetch an instruction with an address stored in the program counter. Thus, fetched instructions may be identified as belonging to an execution entity since the microsequencer already knows (from the program counter) the address of the next instruction.

Thus, for example, given a system with several threads, two or more independent
15 program counters may be used. For example, three program counters PC0, PC1, and PC2 may be associated with microsequencers 430, 432, and 434, respectively. The operating system may load the initial states of the program counters PC0, PC1, and PC2 so that the program counters may fetch instructions associated with the different threads. The combination of PC0 and microsequencer 430 keeps track of the program sequence for a first
20 thread, the combination of PC1 and microsequencer 432 keeps track of the program sequence for a second thread, and so forth. When an instruction pointed to by PC0 is fetched, the CPU 200 knows that the instruction belongs to a first thread having, for example, EID 0. The EID is then attached to the instruction in the CPU 200 and to subsequently decoded micro-operations that are stored in the instruction queue 412 for execution by the microsequencer
25 430, 432, or 434 in the sequencer 414.

As shown in Fig. 4, the output port of the sequencer 414 is provided to a pipeline back-end block 415 that includes various functional units, such as for example, an early branch execution unit 416, a fast decoder unit 418, an arithmetic/logic unit (ALU) 420, and an address generator unit (AGU) 422. During execution of one or more micro-operations by
30 the sequencer 414, these functional units may be accessed to perform requested functions.

The pipeline back-end block 415 also includes register files 424, 426, and 428. The register files 424, 426 and 428 in the CPU 200 correspond to the three EID groups EID 0,

- 16 -

EID 1, and EID 2. The register files 424, 426 and 428 may each include control registers, status registers, flag registers, and general purpose registers. The register files 424, 426, and 428 are updated by the functional units in the pipeline back-end block 415 during operation. According to one embodiment, the register files 424, 426, and 428 may also be accessible
5 independently and concurrently.

In the illustrated embodiment, requests associated with different EIDs may be processed concurrently provided that there are no dependencies among the requests, and further, the multiple requests do not need to utilize the same functional units 416, 418, 420, and 422. During concurrent operation of the microsequencers 430, 432, and 434, the register
10 files 424, 426, and 428 as well as the cache modules in the multi-unit cache memory may be accessed and updated concurrently.

In the data path of the CPU 200, a store buffer 450 (for write operations) and a load buffer 452 (for read operations) store data that are retrieved from or targeted for the BFU 404. The store and load buffers 450 and 452 are coupled to an internal data bus 454 that is
15 coupled to several units, including the P0-cache 102, the P1-cache 104, the P2-cache 106, the pipeline back-end block 415, and a translation look aside buffer (TLB) 456.

Addresses of instructions in the instruction cache 406 are fed to the TLB 456, which is basically a high speed memory in the CPU 200 that translates the virtual address from the instruction cache 406 into a physical address to access the data cache modules 102, 104 and
20 106.

Based on the multi-unit data cache availability, the control logic 436 in the microcode sequencer 414 may select an appropriate instruction for processing by one of the microsequencers 430, 432, and 434. If data access is needed, the microsequencers 430, 432, and 434 may concurrently access the several modules in the multi-unit data cache.

25 Thus, to improve system performance, multiple instructions may be executed in the CPU 200 with concurrent access to data in the L1 multi-unit cache 202.

In some embodiments, the control logic 436 of the sequencer 414 may also consider possible load/store ordering, outstanding data cache refilling, and other issues. For example, in one embodiment, instructions associated with a request that has been determined to have a
30 high hit ratio may be scheduled first, as may instructions of a real-time constrained execution entity having high priority.

- 17 -

Referring to Fig. 5, the general flow of an instruction execution sequence according to an embodiment is illustrated. Instructions are fetched from main memory 206 or L2 cache 204 (at 302) by the CPU 200 over the host bus 203. In one embodiment, associated EIDs are retrieved with the instructions. In another embodiment, the associated EIDs are not stored and thus are not retrieved. The fetched instructions are then translated into internal micro-operations (at 304) by the decoder stage 408, with a corresponding EID attached to each micro-operation. The EID may be the one fetched with the instruction or it may be determined by the CPU 200 based on the address location of the instruction or which microsequencer fetched the instruction. Next, the translated micro-operation is stored in the instruction queue 412 (at 306). The micro-operation is then delivered to one of the microsequencers 430, 432, and 434 for execution (at 308). Execution of the micro-operation may cause a data cache access request to be made (at 310), in which case a corresponding one of the P-cache modules is accessed based on the attached EID. The EID is decoded by the cache controller 108 and an appropriate request is sent to a corresponding P-cache (102, 104, or 106). The data access request is then completed in the assigned P-cache (at 312).

While the embodiments described include a multi-unit cache memory to store data, it is contemplated that the multi-unit cache memory may be adapted to store instructions of different execution entities in further embodiments. In such embodiments, the information stored in the multi-unit cache memory includes the instructions themselves.

While the invention has been disclosed with respect to a limited number of embodiments, those skilled in the art will appreciate numerous modifications and variations therefrom. It is intended that the appended claims cover all such modifications and variations as fall within the true spirit and scope of the invention.

What is claimed is:

- 18 -

- 1 1. A system comprising:
2 a processor;
3 a plurality of execution entities executable on the processor; and
4 a buffer memory having multiple buffer sections, each buffer section adapted
5 to store information associated with requests from a corresponding one of the multiple
6 execution entities.
- 1 2. The system of claim 1, further comprising a software layer adapted to assign
2 identifiers to the execution identifiers and to assign each buffer section an identifier.
- 1 3. The system of claim 2, further comprising a controller operatively coupled to
2 the buffer memory to select one of the buffer sections based on an identifier associated with a
3 request from an execution entity.
- 1 4. The system of claim 1, wherein the execution entities include processes.
- 1 5. The system of claim 1, wherein the execution entities include threads.
- 1 6. The system of claim 1, wherein the buffer memory includes a cache memory
2 having multiple cache sections.
- 1 7. The system of claim 6, wherein the cache sections are configurable to have
2 different attributes.
- 1 8. The system of claim 7, wherein the attributes include cache line replacement
2 policies.
- 1 9. The system of claim 7, wherein the attributes include cache update policies.
- 1 10. The system of claim 7, wherein the attributes include cache organization.

- 19 -

1 11. The system of claim 6, wherein the cache memory includes a multi-level
2 cache memory in which at least one level includes a multi-unit cache memory having
3 multiple cache sections.

1 12. The system of claim 1, wherein the execution entities process data according
2 to different temporal and spatial locality characteristics, and each buffer section is configured
3 based on the temporal and spatial locality characteristics.

1 13. A cache memory for use in a system having multiple execution entities, the
2 cache memory comprising:
3 a cache controller; and
4 a plurality of cache sections, the cache controller adapted to store information
5 in one of the cache sections based on which execution entity the information is associated
6 with.

1 14. The cache memory of claim 13, wherein the cache controller includes storage
2 elements programmable with identifiers to identify the one or more execution entities that
3 each cache section is associated with.

1 15. A method of setting up a cache memory having multiple cache sections, the
2 cache memory located in a system having multiple execution entities, the method comprising:
3 assigning an identifier to each execution entity based on which of the multiple
4 cache sections is to be used for the execution entity; and
5 configuring each of the cache sections based on the cache usage characteristics
6 of the one or more execution entities assigned to the cache section.

1 16. The method of claim 15, wherein the configuring includes setting attributes of
2 each cache section.

1 17. The method of claim 16, wherein the attribute setting includes setting a
2 replacement policy for each cache section.

- 20 -

1 18. The method of claim 16, wherein the attribute setting includes setting an
2 update policy for each cache section.

1 19. A memory subsystem in a system having multiple execution entities that
2 generate instructions, comprising:
3 a controller; and
4 a multi-unit buffer memory having multiple buffers,
5 the controller adapted to select one of the buffers to store information
6 associated with an instruction based on which execution entity generated the instruction.

1 20. The memory subsystem of claim 19, wherein the controller includes storage
2 elements corresponding to each of the buffers that are programmable to values identifying the
3 execution entities.

1/5




FIG. 1

2/5





FIG. 2

3/5**FIG. 3**

4/5

203.






FIG. 4

5/5

**FIG. 5****FIG. 6**

INTERNATIONAL SEARCH REPORT

International Application No

PCT/US 99/21249

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 G06F12/08

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	EP 0 856 797 A (SGS THOMSON MICROELECTRONICS) 5 August 1998 (1998-08-05) page 3, line 54 -page 6, line 50 page 7, line 44 - line 54 page 8, line 10 - line 16; claims 1-4,8; figures 1-4 ---	1-10, 12-20
X	DONGWOOK K ET AL: "A PARTITIONED ON-CHIP VIRTUAL CACHE FOR FAST PROCESSORS" JOURNAL OF SYSTEMS ARCHITECTURE,NL,ELSEVIER SCIENCE PUBLISHERS BV., AMSTERDAM, vol. 43, no. 8, May 1997 (1997-05), page 519-531 XP000685730 ISSN: 1383-7621 page 522, left-hand column, paragraph 3 -page 524, right-hand column, line 3 -----	1-4, 6-10, 13-20

☐ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search

19 January 2000

Date of mailing of the international search report

26/01/2000

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Ledrut, P

INTERNATIONAL SEARCH REPORT

Information on patent family members

In tional Application No

PCT/US 99/21249

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
EP 0856797 A	05-08-1998	EP 0856798 A	05-08-1998
		EP 0890149 A	13-01-1999
		WO 9834172 A	06-08-1998
		JP 10232834 A	02-09-1998
		JP 10232839 A	02-09-1998
		JP 11509356 T	17-08-1999
<hr/>			